



June 12, 2023

Alan Davidson  
Administrator  
National Telecommunications and Information Administration  
U.S. Department of Commerce  
1401 Constitution Avenue NW, Room 4725  
Washington, DC 20230

Submitted via [www.regulations.gov](https://www.regulations.gov)

**RE: AI Accountability Policy Request for Comment (NTIA– 2023–0005)**

Dear Administrator Davidson:

The Blue Cross Blue Shield Association (BCBSA) is pleased to have the opportunity to share our comments on the Artificial Intelligence (AI) Accountability Policy Request for Comment (the “RFC”).

BCBSA is a national federation of 34 independent, community-based and locally operated Blue Cross and Blue Shield (BCBS) companies (Plans) that collectively provide health care coverage for one in three Americans. For more than 90 years, BCBS Plans have offered quality health care coverage in all markets across America – serving those who purchase coverage on their own as well as those who obtain coverage through an employer, Medicare and Medicaid.

BCBSA believes that everyone should have access to high-quality health care. It is important to continue research and development of best practices and standards that address algorithm documentation, testing, use and auditing, as well as stakeholder education, in order to improve algorithm transparency, reliability and trustworthiness, while mitigating the potential for unintended consequences, such as adverse bias and inaccuracies. Earning – and maintaining – consumer trust is of paramount importance, and it can be achieved by efforts on the part of NTIA, other federal partners and the industry to develop self-regulatory, regulatory and other measures and policies that further advance trustworthy AI. We commend NTIA for engaging on this topic.

BCBSA's commitment to the health of our communities also includes continuing to improve the way we gain insight from diverse health factors and how we use technologies. These efforts will improve the delivery of patient-focused care programs, including providing doctors and caregivers the tools they need to improve their care practices and empower consumer choice with in-depth quality and cost information on their care. BCBS Plans are actively leveraging technology, where appropriate, to provide innovative solutions and services to members.

Informed by our experience, BCBSA respectfully offers comments to the RFC. We wish to highlight the following considerations that we believe are particularly important for NTIA to consider as it drafts and issues its report on AI accountability policy development:

- **Alignment with NIST AI Risk Management Framework:** We applaud NTIA for aligning in many ways in this RFC with the approach reflected in NIST's AI Risk Management Framework (RMF) regarding the attributes of trustworthy AI and the acknowledgement that there is no one-size-fits-all approach. As both the RFC and AI RMF reflect, a workable framework is one that is driven by industry input and practice and is use-case specific. We provide additional feedback to reinforce alignment with the AI RMF, where doing so is particularly important.
- **Alignment on AI Regulation.** To the extent regulators take a regulation-based approach as opposed to a self-regulation or voluntary approach in deploying AI accountability measures, we urge NTIA to encourage its federal regulatory agency partners to align on any enforceable requirements in this domain. Regulatory agencies should be advised to only adopt requirements after gathering a deep understanding of the existing regulatory landscape, developments already underway, and impacts (costs and benefits) to consumers.

We appreciate your consideration of our comments and believe that our recommendations will help with the development and use of trustworthy AI through sensible public policies. If you have additional questions or comments, please contact Lauren Choi, Managing Director, Health Data and Technology Policy at [lauren.choi@bcbsa.com](mailto:lauren.choi@bcbsa.com).

Sincerely,



Anshu Choudhri  
Vice President, Policy Development and Strategy  
Policy and Advocacy

## **Detailed Comments**

### **Scope of AI Definition and Framework**

The RFC incorporates the definition developed by NIST for “AI system” as “an engineered or machine-based system that can, for a given set of objectives, generate outputs such as predictions, recommendations, or decisions influencing real or virtual environments.” However, the RFC goes on to further add on an additional layer of scope from the language used in the White House Blueprint for an AI Bill of Rights (“White House Blueprint”) which covers “automated systems” with “the potential to meaningfully impact the American public’s rights, opportunities, or access to critical resources or services.” We ask for NTIA to clarify the scope of the AI definition, particularly as it relates to algorithms that may predict physical or mental health risks. Ideally, the definition adopted by NTIA would be a singular and cohesive one that is adopted based on industry feedback. We note that unlike the NIST Risk Management Framework (RMF) which was produced pursuant to congressional directives and involved several rounds of stakeholder input, the White House Blueprint was not developed with the same robustness of stakeholder input. We believe the NIST definition is sound, with the additional clarification as mentioned.

### **Specific Responses to RFC Questions**

#### **Issue: AI Accountability Objectives**

1. The RFC discusses the purpose of AI accountability mechanisms such as certifications, audits, and assessments.

**Question 1c:** An audit or assessment may be used to verify a claim, verify compliance with legal standards, or assure compliance with non-binding trustworthy AI goals. Do these differences impact how audits or assessments are structured, credentialed, or communicated?

**Response:** BCBSA respectfully suggests that NTIA steer clear of recommending that an external or third-party audit be used to verify compliance with legal standards and focus at this time on deploying external audits, where feasible, as one mechanism to assure consistency with non-binding trustworthy AI goals. Third-party audits are immature as a mechanism to detect or mitigate adverse bias; rather as it stands today, third-party audits are resource-intensive, difficult to obtain, and often cost-prohibitive for businesses when they can be secured. Given that standards for conducting audits are not yet defined, attention and resources should be paid at this time to developing such standards rather than imposing inadequate standards that might fail to identify issues, giving a false sense of security regarding the trustworthiness of the AI system.

**Question 1e:** Can AI accountability practices have meaningful impact in the absence of legal standards and enforceable risk thresholds? What is the role for courts, legislatures, and rulemaking bodies?

**Response:** AI accountability practices, particularly impact assessments, can and do have meaningful impact when they are developed with each industry and relevant use-cases. BCBSA cautions, however, against one-size fits all rulemaking around the use of AI, as its use and application to differing industries are evolving at a record pace. Instead, BCBSA recommends that NTIA support industry's adoption and use of standardized impact assessments, which will enable awareness, transparency and thoughtful analysis of how AI is being used to support a variety of use cases.

3. **Question:** AI accountability measures have been proposed in connection with many different goals, including the eight goals listed in the RFC. To what extent are there tradeoffs among these goals? To what extent can these inquiries be conducted by a single team or instrument?

**Response:** Regarding the fourth goal describing adequate transparency and explanation to affected people about the uses, capabilities, and limitations of the AI system, we caution NTIA that it may be necessary to balance the level of transparency, especially in cases where PHI or sensitive data may be required for greater model accuracy in healthcare. We suggest regulators allow for some degree of flexibility where a model's more favorable accuracy rate is in the public interest, weighing the risks and benefits to consumers. One mitigation strategy would be to allow for an AI system developer to explain why a system is not fully transparent (e.g., use of a dataset containing PHI or sensitive data) or corrective action plan to bridge the system to a 100% transparent one over time to limit any adverse effects from a clinical standpoint.

The goals laid out in the RFC use certain terms like "safety and effectiveness," "algorithmic discrimination protection," "data privacy," "notice and explanation," and "human alternatives, considerations, and fallbacks" that are not defined. As NIST develops its report, it is important to define these terms with precision to provide sufficient guidance to stakeholders so that the efficacy of accountability mechanisms can be understood. Referencing actual definitions, statutes, or policies that these terms are referencing (by industry or use-case) will be important to provide sufficient clarity.

The complexities inherent in these terms underscores the importance of alignment with ongoing federal work in the field of AI accountability measures to achieve these goals, including NIST's AI RMF and other guidelines in development. These measures should consider use-cases and a diverse array of AI applications.

6. **Question:** The application of accountability measures (whether voluntary or regulatory) is more straightforward for some trustworthy AI goals than for others. With respect to which trustworthy AI goals are there existing requirements or standards? Are there any trustworthy AI goals that are not amenable to requirements or standards? How should accountability policies, whether governmental or non- governmental, treat these differences?

**Response:** We agree that some goals are not amenable to requirements or standards. Specifically, with respect to the first goal that the “AI system does not substantially contribute to harmful discrimination against people,” BCBSA asks NTIA to avoid recommending a mechanism that would purport to prove a nullity—here, the absence of harmful discrimination. Mechanisms should be developed using established criteria for what is harmful discrimination so that the AI system can be tested for these criteria. Otherwise, it would be difficult if not impossible to meet this important goal.

As a general matter, as NIST considers these challenges and how governmental and non-governmental actors should apply accountability measures to meet these goals, we ask NTIA continue to coordinate with NIST and other federal entities regarding AI standards.

7. **Question:** Are there ways in which accountability mechanisms are unlikely to further, and might even frustrate, the development of trustworthy AI? Are there accountability mechanisms that unduly impact AI innovation and the competitiveness of U.S. developers?

**Response:** We believe audits and assessments are examples of mechanisms that might frustrate the development of trustworthy AI and could make American industries less competitive relative to peers if not approached responsibly. To the extent audits or assessments are supported by NTIA, it is important to develop a clear and defined scope for such audits, as well as a cadence and compliance timeline that is appropriate and proportionate to the utility of the audit or assessment as well as the resources required to conduct the audit. Under current conditions, annual audits are not feasible, for example. We reiterate our response above that third-party audits are immature and are resource-intensive, difficult to obtain, and often cost-prohibitive for businesses when they can be secured.

8. **Question:** What are the best definitions of and relationships between AI accountability, assurance, assessments, audits, and other relevant terms?

**Response:** As a general matter, BCBSA recommends NTIA look to adopt definitions that are precise to distinguish these terms from each other when doing so is important. NTIA should not develop definitions out of whole cloth but should instead seek to align with those developed by established entities like standards-setting organizations, federal coordinating agencies like NIST, or where relevant for a particular use-case, a federal agency-developed regulation or guidance.

#### **Issue: Existing Resources and Models**

9. **Question:** What AI accountability mechanisms are currently being used? Are the accountability frameworks of certain sectors, industries, or market participants especially mature as compared to others? Which industry, civil society, or governmental accountability instruments, guidelines, or policies are most appropriate

for implementation and operationalization at scale in the United States? Who are the people currently doing AI accountability work?

**Response:** We suggest that the NIST AI RMF is most appropriate for implementation and operationalization at scale in the United States. The AI RMF and RMF Playbook have been generally well received and have the potential for adoption among different industries if done in coordination with other federal agency partners, stakeholder groups, and industry in keeping with the process followed to date. We encourage NIST continue its collaborative approach to date as it builds standalone applications or profiles for a specific industry using the AI RMF.

11. **Question:** What lessons can be learned from accountability processes and policies in cybersecurity, privacy, finance, or other areas?

**Response:** BCBSA thanks NTIA for considering whether there are other areas where accountability processes and policies could be adopted in the development of trustworthy AI. Regarding privacy, we believe that the NTIA should look to the approach taken under the Health Insurance Portability and Accountability Act (HIPAA) where regulated entities are required to make regular assessments based on the facts and circumstances of their individual organizations, including assessments related to privacy breaches and dissemination of privacy notices. These assessments recognize that impacts vary for any given issue and those impacts and risks need to be assessed on a case-by-case basis. Further, HIPAA holds business associates, or third parties that conduct business on behalf of regulated entities, accountable to protect the privacy of individually identifiable health information based on a reasonableness standard. Regarding cybersecurity, we believe that vendor oversight (or third-party service provider) issues are particularly well addressed by the NAIC's Insurance Data Security Model Law and the New York Department of Financial Services (NY DFS) Cybersecurity Regulation. Regarding finance, one best practice is found in state fintech safe harbor statutes and "sandboxes" which foster innovation while setting appropriate on-ramp parameters for new entrants, particularly those from startups or medium-sized firms.

## **Issue: Accountability Subjects**

16. The RFC discusses the lifecycle of a given AI system or component which presents distinct junctures for assessment, audit, and other measures, such as when bias may be most prevalent. The RFC asks a series of questions about how AI accountability mechanisms should consider the AI lifecycle.

**Question 16a:** Should AI accountability mechanisms focus narrowly on the technical characteristics of a defined model and relevant data? Or should they feature other aspects of the socio-technical system, including the system in which the AI is embedded? When is the narrower scope better and when is the broader better? How can the scope and limitations of the accountability mechanism be effectively communicated to outside stakeholders?

**Response:** BCBSA supports NTIA's thinking on this issue to be in alignment with the NIST AI RMF and White Paper, *Mitigating AI/ML Bias in Context: Establishing Practices for Testing, Evaluation, Verification, and Validation of AI Systems*. In these publications, NIST offers recommended guidance for risk management across the AI lifecycle, including employing a socio-technical approach because an AI system is both a technological advancement and a human creation that intersects with our society at different touchpoints.

Regarding how AI accountability mechanisms might address bias, we agree that appropriate controls should be put into place to mitigate adverse bias through the lifecycle. In some contexts where the AI system can be compared to the existing status quo process, a reference set or comparison of biases approach is warranted. An exercise that compares the AI system vs. the "human" system (the system that exists without AI's application) could be instructive to instill trust in the AI system. For example, if such a comparison reveals that the "human" bias rate as the reference bias rate is actually higher than the AI system's bias rate, there is less bias in utilizing the AI system. This result should be factored in when determining whether and when to use the system.

**Question 16c:** How often should audits or assessments be conducted, and what are the factors that should inform this decision? How can entities operationalize the notion of continuous auditing and communicate the results?

**Response:** BCBSA cautions NTIA with recommending a specific cadence for audits and agrees that the decision to audit should be informed through careful consideration of many factors. The cadence should be industry- and even AI system-specific with compliance timelines that are reasonable and feasible, given the level of effort required to conduct a reliable audit. Absent a rare instance of serious noncompliance warranting a corrective action plan that may require more frequent audits, an annual audit for each AI system would not be possible for organizations of any size to manage. Instead, we recommend NTIA look to HIPAA's material change threshold as providing guidance for the circumstances that may trigger an audit.

17. **Question:** How should AI accountability measures be scoped (whether voluntary or mandatory) depending on the risk of the technology and/or of the deployment context? If so, how should risk be calculated and by whom?

**Response:** BCBSA recommends AI accountability measures be scoped applying an impact-based approach. The higher the impact to individuals from the use of the technology or deployment, the higher the risk and need for greater accountability measures.

## **Issue: Accountability Inputs and Transparency**

20. **Question:** What sorts of records (e.g., logs, versions, model selection, data

selection) and other documentation should developers and deployers of AI systems keep in order to support AI accountability? How long should this documentation be retained? Are there design principles (including technical design) for AI systems that would foster accountability-by-design?

**Response:** First, consistent with existing business practice, records that would be maintained on AI systems to support accountability measures would be included in record and information management policies. These policies would indicate what information is retained and stored for the AI system. We see these policies as a likely starting place to house AI-specific record keeping. We appreciate that these policies would need to be updated periodically and on an as-needed basis to take into account the company's AI systems.

Second, we are concerned that AI accountability measures might over-emphasize the record keeping aspect of controls in a way that becomes unwieldy as AI systems continue to be developed and deployed (for some organizations, tens of thousands of systems might be deployed). It is important for policymakers and advisory bodies like NIST and NTIA to consider the real-world application of this accountability measure in terms of the amount of time, financial resources, and workforce needed to audit or keep a record of logs, models, model selection, governance, and data selection and indices. BCBSA recommends a balance between adopting a reasonable record keeping policy and devoting the resources needed to generate better outcomes and value for the users or beneficiaries of an AI system.

21. **Question:** What are the obstacles to the flow of information necessary for AI accountability either within an organization or to outside examiners? What policies might ease researcher and other third-party access to inputs necessary to conduct AI audits or assessments?

**Response:** In the health care sector, policymakers must consider HIPAA's and the HITECH Act's restrictions on information flow. These authorities will inform the degree to which a researcher or third-party may access protected health information (PHI) and electronic PHI (ePHI) when conducting an AI audit or assessment. This example demonstrates the value of NTIA, NIST and other federal agencies engaging in industry-specific discussions to identify the industry's best practices and practices for documenting approvals for algorithms as they move through phases of development, as well as requests for access to inputs by third parties.

22. **Question:** How should the accountability process address data quality and data voids of different kinds? For example, in the context of automated employment decision tools, there may be no historical data available for assessing the performance of a newly deployed, custom-built tool. For a tool deployed by other firms, there may be data a vendor has access to, but the audited firm itself lacks. In some cases, the vendor itself may have intentionally limited its own data collection and access for privacy and security purposes. How should AI accountability requirements or practices deal with these data issues? What should be the roles of

government, civil society, and academia in providing useful data sets (synthetic or otherwise) to fill gaps and create equitable access to data?

**Response:** The complexities described in the examples in this question underscore the need for policymakers to engage with industry and develop industry-specific AI accountability best practice measures alongside these industry stakeholders. Such best practice measures would be to institute controls to address data quality and data voids, such as documentation, a feasibility analysis report, and checklists for industry-specific considerations.

23. **Question:** How should AI accountability “products” (e.g., audit results) be communicated to different stakeholders? Should there be standardized reporting within a sector and/or across sectors? How should the translational work of communicating AI accountability results to affected people and communities be done and supported?

**Response:** The manner by which the results of AI accountability measures should be communicated to stakeholders will vary by industry. Communication methods to affected people and communities should be developed through industry-specific consultation, leveraging any industry-specific AI accountability best practice measures that may exist in the industry about AI systems or more broadly, another industry-specific tool in practice that affects the community or individual.

#### **Issue: AI Accountability Policies**

30. **Question:** What role should government policy have, if any, in the AI accountability ecosystem?

**Response:** We strongly urge government policy on AI accountability to be aligned across the federal government if policymakers take a regulation-based approach, as opposed to a self-regulation or voluntary approach. It is important for regulators to have a clear understanding of the existing regulatory landscape, developments underway, and impacts to consumers and innovation as requirements are created and enforced. Alignment of requirements, where appropriate, will prevent undue burden and regulatory overlap that may stymy innovation and development and impair the leadership of the U.S. in artificial intelligence.

31. **Question:** What specific activities should government fund to advance a strong AI accountability ecosystem?

**Response:** We strongly urge the U.S. government to fund continued research into AI systems and employ strategies to mitigate unintended consequences including bias to improve data equity. We ask NTIA and its federal partner agencies to support research into real-world examples and application of AI and data science principles to ensure robust algorithm development, deployment, and use. This research will help inform development of industry-specific guidelines for AI system accountability

measures where necessary.

In addition, one of the more frequently encountered limitations in AI system trustworthiness relates to the existing data used by the system and whether certain populations are over- or under-represented in the data class. We need government leadership to facilitate efforts to build more accurate data sets by considering the role of existing clinical trials recruitment efforts inside and outside of government, and consumer privacy and rights protections on data equity and quality of AI systems. To this end, we also ask government to continue efforts ongoing inside and outside government to develop further understanding and clarifications that distinguish adverse bias from beneficial bias.